RiskMetrics Group

Research Monthly

The SSG and my two brains

March 2008

Christopher C. Finger
chris.finger@riskmetrics.com

> *Nevertheless, during the turmoil, all firms were able to obtain adequate liquidity to fund their operations.*

With events as fluid as they are, it is a wonder that anyone is able to write anything. No sooner than I start to react to an important document than there is another market event, and the quotation above paints a two-week old document as sadly dated.

But *Observations on Risk Management Practices during the Recent Market Turbulence* is not dated. The document was prepared by the Senior Supervisors Group (SSG) of financial supervisors from France, Germany, Switzerland, the United Kingdom and the United States. Its objective is to survey risk management practices in large financial institutions, and to infer which practices differentiated those banks that performed best through the most recent market turbulence. In fairness, the document is a result of supervisor conversations with (just) sixteen large banking and securities firms, and it is possible that the quotation above still applies as of this writing (or as of your reading) to that sample.

My goal here is not to summarize the document in its entirety, but rather to react to a number of key points that should serve as marching orders to risk management service providers and their clients. Eventually, we will all agree that the subprime crisis, the credit crunch and any other related crises are over, and we will need to assess what we have learned. But even with market turbulence persisting, it is not too early to start learning from our experiences. This is the spirit of the SSG document, and of this reaction to it.

In the document, the SSG makes a number of lists—of recommendations for supervisors, of key management behaviors, of business lines most impacted by recent events—so I will structure my comments around a list as well. There are three areas of risk management—statistical measures, stress testing and liquidity—for which the SSG document provides plenty to ponder. Somewhat opposite to the normal order of things, I will begin where risk modelers deserve the most criticism: liquidity.

Liquidity risk

If there is both good and bad to be found in statistical measures and stress testing, then that leaves liquidity alone to play the role of ugly. That liquidity is at the center of market turbulence is not surprising. We should acknowledge, though, that liquidity is not the dominant characteristic in every period of market turbulence: while the crisis of late 1998 can certainly be pinned on liquidity, the technology bust and corporate scandals of the early part of this decade cannot.

So where are we with liquidity risk, and why does

©2008 RiskMetrics Group, Inc. All Rights Reserved.

this theme keep rearing its head, along with calls for us to model it better? It is not a problem we have ignored—see Malz (2003) for instance—but it is one that has proved ungainly for any practical solutions. As is often the case, the first task is just defining what the problem is. Malz presents a useful discussion of the meaning of liquidity risk, and distinguishes two problems.

One problem is asset liquidity, that is, for a specific asset, how effectively I can transform the asset to cash, without materially altering the market, and independent of the size of my position and the amount I wish to trade. To a large extent, this is about the marketplace, and we can generically characterize trading in a specific security or asset class as liquid or illiquid based on how actively and effectively the assets trade. But there is also a specific element to asset liquidity, in that for positions large enough, we move beyond being simply price takers, and are exposed to the potential for self-induced price deterioration. The last time we discussed liquidity, in the aftermath of the Amaranth trading losses[1], it was with respect to this aspect of liquidity risk.

For asset liquidity, a number of tools have been proposed over the years, none of which having worked its way into common practice. An attractive approach is to suppose a liquidity horizon for individual holdings. The liquidity horizon represents the number of trading days it would take to unwind a position without ever trading in enough size to adversely impact the market. Assigning different liquidity horizons to different positions allows us to acknowledge that in less liquid assets, we are exposed to potential market volatility for a longer time.

[1] See Finger (2006b).

Plenty of practical issues arise, however. One is how we estimate liquidity horizons. Malz discusses this, suggesting a link with the average daily trading volume. Malz also points out the drawbacks of this approach, including that not all financial instruments have observable volume information, and concedes that a trader's intuition is likely the best measure available.

A second issue is how to make sense of our risk horizon once we have a mix of different liquidity horizons in our portfolio. This inevitably leads us to the theme of portfolio aging, which we raised in our last note. As a start, we can take the crude approach of scaling up the volatility according to the liquidity horizon and running our standard single risk horizon analysis. This is not a perfect solution, but at least addresses the first order problem. More importantly, it pushes us just to ask the question about liquidity horizon and to observe our position sizes in relation to volumes, exercises which on their own raise awareness of liquidity issues even if the information is not incorporated into formal models.

A third issue is the credibility of a risk measure with an assumption that all positions will be liquidated. These days, complaints that measures are overly conservative will fall on deaf ears, but in normal times, a full liquidation risk measure will likely prove so large as to eventually be ignored. It would make sense then to examine VaR under both the standard mark-to-market assumptions as well as the full liquidation framework, thereby illuminating our exposure to potential liquidity crises.

The second problem is funding liquidity: whether an institution's assets and funding arrangements are suf-

ficient to meet their obligations and to continue normal operations. This is for the most part a specific problem, depending on the institution's relationships with its counterparties and creditors. Indeed, with recent events, funding liquidity was the endgame, with Bear Stearns ultimately unable to obtain short-term financing.

In the Bear Stearns case, though, asset liquidity played the role of catalyst. Among U.S. brokers, Bear was in the unfortunate position of being particularly reliant on asset-based funding, with approximately 27% of their liabilities in the form of repos, and at the same time holding a relatively large portion of its assets in mortgage-related and asset-backed securities.[2] This left Bear uniquely exposed to the deterioration in mortgage-backed securities, and to their creditors' subsequent aversion to these as collateral for funding.

With respect to funding liquidity, the SSG highlights the practices of the more successful banks, in particular the alignment of the treasury (funding) operations with risk management and the accurate assessment of internal charges for funding needs, reflecting both normal operations and the need for contingent liquidity in periods of market stress. Moreover, the SSG endorses the practice of asking questions about events or actions that could lead to reduced liquidity.

Funding liquidity being largely a specific, rather than a market, issue, the responsibility for asking these questions falls more on the business managers than the risk modelers. Do multiple sources of liquidity exist? What constraints can I expect, contractually or otherwise, should markets come under stress? Can a counterparty change liquidity provisions, such as what collateral they will accept, with little notice?

These are not just questions for banks and dealers, but for hedge funds as well. Do multiple prime brokerage agreements provide adequate diversity for funding? What would be the impact of more stringent margin requirements?

So there is a strong emphasis on financial institutions examining these questions, ideally during benign periods as well as turbulent ones. From a modeling point of view, it is too much of a stretch to ask for the probability that a prime broker tightens its margin rules. It is reasonable, though, to ask for a modeling framework in which we can examine scenarios of tighter credit, and to assess the impact of a hypothetical change in requirements.

For liquidity in general, that we do not have any perfect models is acceptable. That we do not have any established models, with a history of failures, successes and dialogues is not. As a community, we need to get on with implementing what we have, gaining experience and making things better.

Statistical measures

In any severe market, there will be calls to abandon statistical risk forecasts (lumped typically, and wrongly, under the label of Value-at-Risk) altogether. To its credit, the SSG takes a more measured approach, giving credit where it is due, and pointing out a few practices that lead to bad forecasts. Overall, their discussion of statistical measures comes down to two conclusions: bad models perform badly, and it is important to get the details right.

[2] See Rosenberg (2008).

The SSG begins its discussion of Value-at-Risk practices with the diplomatic statement

> *While most firms reported that their VaR systems generally worked as expected . . . many firms identified weaknesses in their particular implementation of VaR . . .*

The discussion proceeds to reference cases of abnormally frequent backtesting excessions (that is, days where losses exceeded the VaR forecast) and to mention that several banks were considering making their volatility models more dynamic. A bit later, there is a discussion of conditional versus unconditional market risk measures, and the statement that banks desired to move to more conditional measures.

The tone of this discussion implies that there is still widespread use of historical simulations. (To be clear, the SSG does not mention any model explicitly; it is this author's intuition that "unconditional" is a code for historical simulations.) I have written before[3] that while this technique is attractive for its simplicity, its forecasts perform poorly, with much of the blame on its implicit equal weighting of historical data. Its adoption represents too great a sacrifice of performance for the benefit of transparency. This is not an indictment of the entire modeling community, however: better, yet still practical, models do exist, even if they are not always used.[4]

Within the discussion of poor model performance is an important observation. Some firms reported managing risk based not on regulatory guidelines (99% confidence level VaR) but a lower threshold. This had two effects: one was that it made the numbers more relevant to management; second was that it meant the threshold was breached more often, meaning there could be more discussions about whether the model was performing well. If we question our risk forecasts, we can make them better. This is a lesson we should remember in the benign times as well, as a lack of VaR excessions can be just as important an indicator of bad model performance as an abundance of them.

Moving to the details, the SSG highlights two issues—the CDS-bond basis and the use of proxies for illiquid assets—both of which involve the choice of data to model different instruments. If a risk model were a rock-and-roll band, the volatility model and statistical distribution would be the lead singer and lead guitar player. The choice of data would be the bass player—unnoticed when all goes well, appreciated only by the most discerning fans but utterly essential to the end product. The SSG is to be commended for noticing the bass player's importance.

A CDS-bond basis trade comprises a long credit position in the form of a bond and a short credit position in the form of a CDS referencing the bond's issuer. The SSG points out that in numerous cases, such trades were attributed zero risk. Clearly, the risk on such trades should be minimal, in that the primary source of risk—the credit of the issuer—is canceled out; but the two sides of the trade can still vary. Failing to account for this variation leads to an understatement of risk.

In fact, this is a typical problem with any basis trade.

[3] See Finger (2006a).

[4] See Zumbach (2007).

In one sense, a basis arises from the observation that one financial instrument is almost a replacement for another. To the extent that two instruments are in fact used interchangeably (going long credit by buying a bond or selling credit derivative protection, going long interest rates through bonds or swaps), it is sensible to model the two instruments with one source of data. In this way, we capture the essential source of risk without adding needlessly to the number of factors we need to manage. But when trading evolves to taking positions on the basis, that is, exploiting the fact that the two instruments are not perfect substitutes, then our risk model must reflect this complexity and incorporate more factors. In short, we can ignore the basis until we start trading it.[5]

The second technical issue is the use of AAA-rated corporate bonds as a proxy for the risk of similarly rated subprime-backed CDOs. Clearly, this practice caused substantial understatements of risk. Some will point to the rating agencies as having misrepresented their ratings as indicative of similar credit quality across structured finance and corporate debt, but this is beside the point. Even if the ratings did indicate similar default risk—which is the most any of the rating agencies would have claimed—the proxy would still not have been defensible.

First, the use of the proxy was for price risk, which the rating agencies do not claim to represent. Second, beyond price risk, the risk proxy also is intended to provide for the relationship between the CDOs and other market factors. Here, the proxy is hard to accept in general, but presents particular problems for the relationship between CDOs and the assets that underly them. The use of the proxy means that the relationship between subprime ABS and CDOs that contain them will be through a statistical relationship between the ABS and highly rated corporate bonds. This framework will never succeed in capturing the nonlinearity in highly rated CDOs, that is, that they may be insensitive to their underlyings for small moves, but highly sensitive once losses on the underlyings approach their subordination level.

Of course, it is one thing to tear down a practice, and another to provide an alternative. It would be wonderful to treat subprime CDOs as we do tranches of the corporate credit indices—modeling their sensitivity to both the underlying credits and to an observed market correlation through a specific pricing model—but the illiquidity of the underlying subprime ABS and the complexity of the CDO structures preclude this. As a matter of priority, though, rather than focusing on matching the credit quality of the CDOs, our efforts would be better spent on the relationships between the CDOs and the rest of the market. Thus a different crude proxy could be to model the underlying portfolio of ABS through the best price information available, and then capture the relationship between the ABS and the CDO through a basic approximation of the CDO structure. At very least, such an approach gives us warnings when we ask simple questions such as what happens if the ABS market collapses.

On statistical risk measures, then, the modeling community gets off lighter than on liquidity risk. There are always practical choices to be made, but within the arsenal of models we have, there are at least reasonably good alternatives. Moreover, these alternatives are not just ideas, but are implemented, inter-

[5] Specific to credit risk, I wrote about this issue in Finger (2005).

nally in some banks as well as in RiskMetrics and other third party services. The practices that the SSG most criticizes, then, are not cases where no answer exists, but rather cases where better alternatives were neglected.

Stress testing

The first issues the SSG raises on stress testing overlap those with statistical risk measures. In particular, the bass player makes an appearance, in that if exposures are not linked to appropriate risk factors, then stressing the right factors may still not give us any information. Recalling the CDO proxy, we could have been prescient about ABS deterioration, but our stress tests would not have shown CDO losses if we were modeling CDOs with corporate bond data.

Specific to stress testing, one interesting point is that some institutions reported challenges in getting senior managers to seriously consider the stress tests the risk managers applied. Though part of the blame here could be a lack of senior management involvement in risk, there is also a burden of responsibility that falls on the risk managers.

Stress tests that are no more than large moves in a small number of key risk factors probably deserve to get ignored. Since we accept that the art of stress testing will never illuminate the precise scenario that ultimately plays out, the value in the exercise is in the dialogue that gets us to the scenario. Why is the scenario relevant to our position? What events could cause such a scenario to occur? Would the scenario also affect our creditors or trading partners? Would market participants' reaction to the event produce market contagion? A conversation about these questions is a more beneficial exercise than the simple mechanics of calculating the loss we would incur, and is not something managers can easily ignore.

The SSG's last point is that useful stress tests were informed by the business practices of the firm. This means that risk managers should consider how the firm makes money as well as how it funds itself. Stress tests, informed by this, should identify dangerous market moves, but should also consider how the firm, its competitors and its counterparties might react. One direction this can lead is back to liquidity.

With a leveraged position in a risky security, we are exposed both to funding liquidity (how we obtained the leverage in the first place) and to asset liquidity. A number of events can impact the state of this liquidity. Suppose a competitor (or many) hold a similar position with a different funding counterparty. This counterparty could tighten credit, forcing our competitor to sell his risky position, eroding the asset liquidity of our risky position, prompting our funding counterparty to tighten their terms with us. These are events that are unlikely to surface in any of the data from which we would build statistical models (price histories, trading volumes), and thus it is scenarios like these that stress testing should address.

The most narrow view of stress testing is that it is a means to examine market scenarios that are not considered by our statistical models. In this sense, they serve as a complement to the models, and their mechanics are largely the same as what we do under Monte Carlo: we generate scenarios, by a statistical model on one hand or by expert judgement on the other, and then apply our pricing models to assess our loss. But the best practices identified by the SSG

go beyond this, delving into the business practices of the firm and considering not simply a market scenario but our, and others', reaction to it.

Conclusion

We in the risk community are fond of saying that risk management is as much about art as about science. While true, this statement is dangerous, in that it tempts us to use it to dismiss failings of the scientific part of our jobs. We need to remind ourselves that there is value in statistical models, but that if applied poorly, they give poor results. And we have to admit that our models are not complete, particularly regarding liquidity, though there are ideas floating around that we should begin to gain experience with. The appropriate response to all this is to improve our science, not to dump all the problems to the other side of our collective brain.

Our right brain does not escape either, but, as some of us may have experienced in dealing with the more right-brained of our kind, it is in need of a bit of focus. The job of the art in risk management needs to be more than just coming up with odd market scenarios to pass to our valuation machine, a sop for a disgruntled quant advocating complex but intractable probability distributions.

The job of art needs to be not simply providing different answers, but asking different questions, sparking discussions about business practices, about how we would react in hypothetical situations, about whether we think others would do the same. Our models should facilitate these practices. If there is one crucial lesson, it is that those institutions that built a culture around such discussions have fared the best so far, and will likely continue to do so.

Further reading

- Finger, C. (2005). Spread Values, RiskMetrics Research Monthly, November.
- Finger, C. (2006a). How Historical Simulation Made Me Lazy, RiskMetrics Research Monthly, April.
- Finger, C. (2006b). The Lights Are On, RiskMetrics Research Monthly, October.
- Malz, A. (2003). Liquidity Risk: Current Research and Practice. *RiskMetrics Journal*, **4**(1): 35–72.
- Rosenberg, J. (2008). Situation Room: Bear Market, Bank of America. March 14.
- Senior Supervisors Group (2008). *Observations on Risk Management Practices during the Recent Market Turbulence*. March 6.
- Zumbach, G. (2007). Backtesting Risk Methodologies from One Day to One Year. *RiskMetrics Journal*, **7**(1): 17–60.